

ABSTRACT OF THE DISCLOSURE

An architecture and method are presented for a computer processor supporting interleaved execution of multiple concurrently-active threads, and capable of independently allocating a portion of the total processor execution time to each of the threads. Compared to existing architectures, in which the portion of processor time allocated to each thread is fixed, the processor architecture described herein is believed to offer higher performance for applications such as communications protocol processing, in which the workload of individual threads may vary, and in which the workload requires real time facilities.